arXiv:1607.07598v2 [math.OC] 19 Aug 2016

Submodular Search is Scheduling

Robbert Fokkink
Department of Applied Mathematics
TU Delft
Mekelweg 4
2628 CD Delft
Netherlands
r.j.fokkink@tudelft.nl

Thomas Lidbetter
Department of Mathematics
London School of Economics
London WC2A 2AE
United Kingdom
t.r.lidbetter@lse.ac.uk

László A. Végh
Department of Mathematics
London School of Economics
London WC2A 2AE
United Kingdom
l.vegh@lse.ac.uk

August 22, 2016

Abstract

Suppose an object is hidden according to a known probability distribution in a finite set S of hiding places which must be examined in some order. The cost of searching subsets of S is given by a submodular function and we seek an ordering of S that finds the object in minimal expected cost. This problem is NP-hard and generalizes known search problems as well as the scheduling problem $1|prec|\sum w_j g(C_j)$ where a set of jobs with weights and precedence constraints must be ordered to minimize the weighted sum of some concave function g of the completion times. We give an efficient combinatorial 2-approximation algorithm for the problem, generalizing analogous results in scheduling theory. We also give better approximations for cost functions with low *total curvature* and we give a full solution for cost functions we call *fully reducible*. Next we suppose the object is hidden adversarily in S and thus we consider the zero-sum game between a cost-maximizing Hider and a cost-minimizing Searcher. We prove that the optimal mixed strategies for the Hider are in the base polyhedron of the cost function, and we solve the game for fully reducible cost functions, giving approximately optimal strategies in other cases.

1 Introduction

We start by considering a search problem with a finite, non-empty set S of hiding locations and a cost function $f\colon 2^S \to [0,\infty)$. An ordering, or *search* π of S must be chosen. For a given ordering π and an element j of S, we denote by $S_j = S_j^\pi$ the union of j and all the locations that precede j in the ordering π. The *search cost* of j under π is $c_f(\pi, j) = f(S_j^\pi)$. We will usually suppress the subscript in c_f unless there is any ambiguity. We assume that a Hider is placed in one of the locations according to a known probability distribution, given by a probability vector $\mathbf{x} \in [0,1]^S$. We seek the ordering π that minimizes the expected search cost with respect to $\mathbf{x}$, which we write as $c(\pi, \mathbf{x}) = \sum_{j=1}^n \mathbf{x}(j)c(\pi, j)$.

We restrict attention to cases when f is submodular and non-decreasing, in which case we call this problem the *submodular search problem*, and if π minimizes $c(\pi, \mathbf{x})$, we say π is *optimal*. The problem is well defined for any arbitrary cost function f, but it has a particularly nice structure when f is submodular, which is a natural assumption for modeling problems with a decreasing marginal cost of searching. The submodular search problem is NP-hard, as we explain later, and our main result, proved in Section 2 is an algorithm computable in strongly polynomial time which approximates the optimal search by a factor of 2. We also show that the algorithm has approximation ratio $1/(1-\kappa)$, where $\kappa \in [0,1]$ is the *total curvature* of f (roughly speaking, the extent to which f differs from a modular function), so that it performs well for cost functions of low total curvature. Later we introduce the concept of a *fully reducible* submodular function, and show how to find an optimal search if f is fully reducible.

If the hiding distribution $\mathbf{x}$ is unknown, then we consider the problem of finding a randomized choice of searches that minimizes the expected search cost in the worst case. This is equivalent to solving a finite zero-sum game between a Searcher and a cost maximizing Hider. A pure strategy for the Searcher is a permutation π of S and a pure strategy for the Hider is an element $i \in S$; the payoff is $c(\pi, i)$. We call this the *submodular search game*. Since the strategy sets are finite, the game has a value and optimal mixed strategies. However the size of the Searcher's strategy set makes the problem of computing these optimal strategies hard. This is a *search game with an immobile Hider in discrete locations*, which is a type of game that has been well studied, see [1, 15, 16]. It is customary to study such games on graphs, and the cost is given by the time taken for the Searcher to reach the Hider's location from a giving starting point. The alternative approach in our paper is to ignore the graph and focus on the cost function.

We analyze the submodular search game in Section 3, showing that the optimal Hider strategy lies in the *base polyhedron* of f, and that any such strategy approximates the optimal strategy by a factor of 2. We go on to give $\left(\frac{1}{1-\kappa}\right)$-approximate strategies, where κ is the total curvature of f. Finally, we give a solution of the submodular search game in the case that f is fully reducible.

This model of search was introduced in [13], and the submodular search problem generalizes search problems with one or more hidden object located on a graph, as studied in [3] and

[20]. The problem is also related to single machine scheduling, which is another subject that has been well studied (see [19], for example). Here a set of jobs with precedence constraints must be processed by the machine in some order given by a ordering π . Each job j has a processing time p_j and the cost C_j of job j is equal to the sum of its own processing time and the processing times of all jobs that precede it in π. Each job j also has a weight w_j and the objective is to minimize $\sum_{j\in S} w_j g(C_j)$ for some monotonically increasing function g. This problem may be denoted by $1|prec|\sum w_j g(C_j)$, to adopt standard scheduling notation. For arbitrary functions g nothing is known about the problem. Indeed, without any restrictions on g, it is a difficult to believe anything can be said in general.

In the case that g is the identity, this is the classical problem $1|prec|\sum w_j C_j$ of minimizing the weighted completion time of scheduling precedence constrained jobs on a single machine. It is well known to be NP-hard [17, 19] and there are many 2-approximation algorithms [5, 9, 10, 18, 22, 27, 29]. Improving the factor of 2 is considered one of the most vexing open questions in scheduling [31]. It is also known that there is no polynomial time approximation scheme for the problem unless NP-complete problems can be solved in randomized subexponential time [7]. In the case that g is something other than the identity, very little is known. If there are no precedence constraints and $g(C_j) = C_j^\beta$, this is the problem $1||\sum C_j^\beta$, as studied in [8], in which it is shown that the problem of minimizing total weighted completion time plus total energy requirement (see [12, 24]) can be reduced to $1||\sum C_j^\beta$.

IIn the case that g is concave, such as when $g(C_j) = C_j^\beta$ for $\beta \in [0, 1]$, the problem $1|prec|\sum w_j g(C_j)$ can be interpreted in the framework of the submodular search problem. The concavity of the function g corresponds to the machine benefiting from a learning effect or from a continuous upgrade of its resources. Let S be the set of jobs and let $\mathbf{x}(j) = w_j$. For a subset A of jobs, let $\tilde{A}$ be the precedence closure of A and let $f(A) = g(p(\tilde{A}))$, where $p(\tilde{A})$ is the total processing time of all the jobs in $\tilde{A}$. Then f is submodular, and the problem of finding the optimal schedule is equivalent to the submodular search problem. This implies that the submodular search problem is NP-hard.

Our main result in this paper implies that there is a 2-approximate algorithm for $1|prec|\sum w_j g(C_j)$ executable in strongly polynomial time, when g is concave and non-decreasing. This result is also obtained in a different way in [30].

2 The Submodular Search Problem

Let $S = \{1, \ldots, n\}$ be a finite set. A function $f\colon 2^S \to \mathbb{R}$ is *submodular* if

$$f(A \cup B) + f(A \cap B) \leq f(A) + f(B)$$

for all sets $A, B \subseteq S$.

We consider the submodular search problem, defined in Section 1, with non-decreasing, non-negative submodular cost function f. The optimal search remains the same if we add a

constant to f, and submodularity is preserved, so we may assume that $f(\emptyset) = 0$ (in other words, f is a polymatroid set function). Rescaling by a positive constant also preserves the optimal search and submodularity, so we may assume that $f(S) = 1$. We also assume that $f(A) > 0$ for all $A \neq \emptyset$, since it is clear that sets with zero cost must be searched first.

We call a vector $\mathbf{x} \in [0,1]^S$ whose coordinates are non-negative and add up to one a *Hider distribution* on S. For $A \subset S$, we write $\mathbf{x}(A) = \sum_{i \in A} \mathbf{x}(i)$.

A key concept we will use in the paper is that of the *search density* (or simply *density*) of a set $A \subset S$, which is defined as the ratio of the probability the Hider is located in A by the cost of searching A. Search density is a concept that often appears in the theory of search games (see [2, 3, 4]), and a general principal that arises is that it is best to search regions of higher density first. The corresponding inverse ratio of the processing time to the weight of jobs also arises naturally in scheduling theory, particularly in the well-known Smith's rule [33] for minimizing the weighted completion time in single machine scheduling of jobs without precedence constraints. The rule says that the jobs should be executed in non-decreasing order of this ratio. Our 2-approximation for the submodular search problem relies on a key result that there is an optimal search that begins with a maximum density subset of S. The proof of this result and the resulting 2-approximation is inspired by the analogous result in [9], of which this is a generalization.

Definition 1. *Let* $\mathbf{x}$ *be a Hider distribution on* S*. The* search density *(or simply* density*) of a non-empty subset* $A \subset S$ *is defined as*

$$\rho(A) = \frac{\mathbf{x}(A)}{f(A)}.$$

We denote $\max\{\rho(A) \colon A \subseteq S\}$ *by* ρ^* *and if* $\rho(A) = \rho^*$ *then we say that* A *has* maximum search density*, or simply* maximum density*. We put* $\rho(\emptyset) = \rho^*$*.*

Recall that $\mathcal{F} \subset 2^S$ is a *lattice* if $A, B \in \mathcal{F}$ implies that $A \cup B \in \mathcal{F}$ and $A \cap B \in \mathcal{F}$. A non-empty $A \in \mathcal{F}$ is an *atom* if the only proper subset of A in $\mathcal{F}$ is the empty set. Atoms are disjoint and each element of $\mathcal{F}$ is a union of atoms.

If f_1 and f_2 are submodular functions on sets S_1 and S_2 then the *direct sum* $f_1 \oplus f_2$ of f_1 and f_2 over S_1 and S_2 is the set function on $S_1 \cup S_2$ defined by

$$(f_1 \oplus f_2)(A) = f_1(S_1 \cap A) + f_2(S_2 \cap A).$$

The restriction of f to a subset B is called the *reduction* f^B. We say that a proper non-empty subset $B \subset S$ is a *separator* if f is the direct sum of f^B and $f^{\bar{B}}$, where $\bar{B}$ denotes the complement of B. In order to check that B is a separator, we only need to verify that $f(S) = f(B) + f(\bar{B})$ (see [11, Proposition 5])

In the proof of Lemma 1, and later in the proof of Lemma 6, we use the observation that $\frac{a}{b} \leq \frac{c}{d}$ implies $\frac{a}{b} \leq \frac{a+c}{b+d} \leq \frac{c}{d}$ for non-negative real numbers a, b, c, d with $b, d \neq 0$. Furthermore, if one of these three inequalities is an equality, then all the inequalities are equalities.

Lemma 1. *Let $\mathcal{M}$ be the family of subsets of maximum density and let M be the union of all the atoms of $\mathcal{M}$. Then $\mathcal{M}$ is a lattice and the cost function f^M is a direct sum over the atoms.*

Proof. If $A, B \in \mathcal{M}$ then $\rho^* = \mathbf{x}(A)/f(A) = \mathbf{x}(B)/f(B)$ and

$$\rho^* = \frac{\mathbf{x}(A) + \mathbf{x}(B)}{f(A) + f(B)} = \frac{\mathbf{x}(A \cup B) + \mathbf{x}(A \cap B)}{f(A) + f(B)} \leq \frac{\mathbf{x}(A \cup B) + \mathbf{x}(A \cap B)}{f(A \cup B) + f(A \cap B)},$$

by submodularity. This inequality is in fact an equality, since $\rho(A \cup B)$ and $\rho(A \cap B)$ are both bounded above by ρ^*. It follows that both $A \cup B$ and $A \cap B$ have maximum density. If A and B are atoms then $A \cap B = \emptyset$, and the equality implies that $f(A) + f(B) = f(A \cup B)$, so $f^{A \cup B}$ is a direct sum over A and B. Therefore, $\mathcal{M}$ is a lattice and f^M is a direct sum over the atoms. □

We say that A is an *initial segment* of a search strategy π if $A = \{\pi(1), \ldots, \pi(|A|)\}$.

Theorem 1. *Let M the largest element of $\mathcal{M}$ and let $\mathbf{x}$ be a Hider distribution. Then any optimal search π has initial segment M. Furthermore, if $A \in \mathcal{M}$, then there exists an optimal search π' such that A is an initial segment.*

Proof. Let A be any subset of maximum search density. Suppose that an optimal search π starts by searching sets $B_1, A_1, B_2, A_2, \ldots, B_k, A_k \subset S$ in that order before searching the rest of S, where $A_i \subset A$ and $B_i \subset \bar{A}$ for all i, and B_1 may be the empty set. Let $A^j = A_1 \cup \ldots \cup A_j$ and similarly for B^j.

Define a new search π' which starts by searching $A_1, \ldots, A_k, B_1, \ldots, B_k$ before searching the rest of S in the same order. Within each A_i and B_i the new search follows the same order as π. Let Δ be the difference between the expected search cost of π and π'. We will show that $\Delta \geq 0$.

First suppose $s \in A_j$ and let C be the subset of A_j containing s and all elements searched before s. Then the difference in the search cost of s under π and π' is

$$f(A^{j-1} \cup B^j \cup C) - f(A^{j-1} \cup C) \geq f(A \cup B^j) - f(A),$$

by submodularity. Now suppose that $t \in B_j$ and let D be the subset of B_j containing t and all elements searched before t. Then the difference in the search cost of t under π and π' is

$$f(A \cup B^{j-1} \cup D) - f(A^{j-1} \cup B^{j-1} \cup D) \leq f(A) - f(A^{j-1}),$$

by submodularity. Hence Δ satisfies

$$\Delta \geq \sum_{j \leq k} \mathbf{x}(A_j) \left[f(A \cup B^j) - f(A) \right] - \sum_{j \leq k} \mathbf{x}(B_j) \left[f(A) - f(A^{j-1}) \right]. \tag{1}$$

Now compute

$$f(A \cup B^j) - f(A) = \frac{\mathbf{x}(A \cup B^j)}{\rho(A \cup B^j)} - \frac{\mathbf{x}(A)}{\rho(A)} \geq \frac{1}{\rho^*} \left(\mathbf{x}(A \cup B^j) - \mathbf{x}(A)\right) = \frac{\mathbf{x}(B^j)}{\rho^*}.$$

Similarly

$$f(A) - f(A^{j-1}) = \frac{\mathbf{x}(A)}{\rho(A)} - \frac{\mathbf{x}(A^{j-1})}{\rho(A^{j-1})} \leq \frac{1}{\rho^*} \left(\mathbf{x}(A) - \mathbf{x}(A^{j-1})\right).$$

Substituting these bounds into (1), we obtain

$$\begin{aligned} \rho^* \Delta &\geq \sum_{j \leq k} \mathbf{x}(A_j)\mathbf{x}(B^j) - \sum_{j \leq k} \mathbf{x}(B_j) \left(\mathbf{x}(A) - \mathbf{x}(A^{j-1})\right) \\ &= \sum_{j \leq k} \mathbf{x}(A_j) \sum_{i \leq j} \mathbf{x}(B_i) - \sum_{j \leq k} \mathbf{x}(B_j) \sum_{i \geq j} \mathbf{x}(A_i) \\ &= 0, \text{ by swapping the order of summation.} \end{aligned}$$

Since π is optimal, it must be true that $\Delta = 0$ and all inequalities above are equalities. It follows that $\rho(A \cup B^j) = \rho(A) = \rho^*$ for all j, and in particular $\rho(A \cup B) = \rho(A \cup B^k) = \rho^*$. We have thus established that if A has maximum search density, then it is a subset of an initial segment $A \cup B$ of maximum density. Therefore, every optimal strategy π searches M first. We have also established that there exists an optimal search that has A as an initial segment. □

2.1 A 2-approximation

Theorem 1 suggests an approach to constructing an optimal strategy. First find a non-empty subset $A \subset S$ of maximum density. By Theorem 1 there is an optimal strategy that begins with the elements of A. Now consider the subproblem of finding an optimal search of $\bar{A}$ with cost function f_A defined for $B \subseteq \bar{A}$ by $f_A(B) = f(A \cup B) - f(A)$ against the Hider strategy $\mathbf{x}$ restricted to $\bar{A}$. The function f_A is called the *contraction* of f by A and is well known to be submodular [14, page 45]. It is easy to see that a search of S that begins with the elements of A is optimal only if it defines an optimal search of $\bar{A}$ with cost function f_A. We summarize the observation below.

Lemma 2. *Suppose there is an optimal search of S with intial segment A. Then an optimal search of S can be found by combining an optimal search of A with respect to cost function f^A with an optimal search of $\bar{A}$ with respect to cost function f_A.*

We now repeat the process on $\bar{A}$ with cost function f_A, finding a subset of maximum density, and so on. The result is a partition of S into subsets $A = A_1, A_2, \ldots, A_k$ such that there exists an optimal search strategy that respects the ordering of those subsets. This is a generalization of the notion of a Sidney decomposition for optimal scheduling [32].

We show that in fact, *any* search that respects the ordering of the decomposition $A_1, \ldots, A_k$ described above approximates the optimal search by a factor of 2, generalizing the analogous result for scheduling that can be found in [9]. We first show that if S itself has maximum density then any search approximates the optimal search by a factor of 2.

Lemma 3. *Suppose that S has maximum search density. Then every search strategy has an expected cost in between $1/2$ and 1.*

Proof. We normalize the cost function to $f(S) = 1$, as usual so that $\rho^* = 1$. Without loss of generality suppose that the search given by $\pi(j) = j$ is optimal so that $S_j = S_j^\pi = \{1, \ldots, j\}$. Then the expected cost of π is

$$\begin{aligned} c(\pi, \mathbf{x}) &= \sum_j \mathbf{x}(j) f(S_j) \\ &\geq \sum_j \mathbf{x}(j)\mathbf{x}(S_j) \text{ (since } \rho(S_j) \leq \rho^*) \\ &= \sum_j \mathbf{x}(j)^2 + \sum_{i<j} \mathbf{x}(i)\mathbf{x}(j) \\ &= \frac{1}{2}\left(\sum_j \mathbf{x}(j)\right)^2 + \frac{1}{2}\sum_j \mathbf{x}(j)^2 \\ &= \frac{1}{2} + \frac{1}{2}\sum_j \mathbf{x}(j)^2 \qquad (2) \\ &\geq \frac{1}{2}. \end{aligned}$$

The maximum expected cost of any search is $f(S) = 1$. It follows that if S has maximum search density then $1/2 \leq c(S) \leq 1$. □

Our 2-approximation relies on being able to find a maximum density subset efficiently. For this, we use a result of Nagano [25] on submodular function optimization, which we state below. First we recall [14, page 34] the definition of the *submodular polyhedron* $\mathbb{P}(f)$ of a submodular function $f : S \to \mathbb{R}$,

$$\mathbb{P}(f) = \{\mathbf{y} \in \mathbb{R}^S : \mathbf{y}(A) \leq f(A), \forall A \subset S\}.$$

Theorem 2 (Theorem 5.1 of [25])**.** *Let $f : 2^V \to \mathbb{R}$ be a submodular function with $f(\emptyset) = 0$ and a starting point $\mathbf{y}_0 \in \mathbb{P}(f)$ and a direction vector $\mathbf{a} \in \mathbb{R}^V$. Then there is an algorithm to solve the following problem that requires a strongly polynomial number of oracle calls.*

$$\textit{Find } t^* = \max\{t \in \mathbb{R} : \mathbf{y}_0 + t\mathbf{a} \in \mathbb{P}(f)\}.$$

Theorem 3. *Suppose that the submodular function f is given by a value giving oracle. Then there is a 2-approximation for the optimal search strategy against $\mathbf{x}$ that can be computed in strongly polynomial time.*

Proof. We first show that a maximum density subset of S can be computed in strongly polynomial time. Applying Theorem 2 with $\mathbf{y}_0 = 0$ and $\mathbf{a} = \mathbf{x}$, we have

$$t^* = \max\{t : \rho(A) \leq 1/t, \forall A \subset S\} = (\min\{s : \rho(A) \leq s, \forall A \subset S\})^{-1} = 1/\rho^*.$$

Hence ρ^* can be computed in strongly polynomial time. A subset A of density ρ^* can then be computed by finding the minimizer of the submodular function $f(A) - \mathbf{x}(A)/\rho^*$. If A is the entire set, then any search is a 2-approximation by Lemma 3. If A is a proper subset, then there exists an optimal search with initial segment A. Let $\mathbf{x}|_A$ be the vector $\mathbf{x}$ restricted to the set A and let $c_f(\mathbf{x}) = \min_\pi c_f(\pi, \mathbf{x})$. Then $c_f(\mathbf{x})$ satisfies

$$c_f(\mathbf{x}) = \mathbf{x}(A)c_{f_A}(\hat{\mathbf{x}}|_A) + \mathbf{x}(\bar{A})f(A) + \mathbf{x}(\bar{A})c_{f^A}(\hat{\mathbf{x}}|_{\bar{A}}),$$

where $\hat{\mathbf{y}} := \mathbf{y}/\|\mathbf{y}\|$ denotes the normalization of a vector $\mathbf{y}$. By induction, if we have a 2-approximation for A and $\bar{A}$ then we have one for S. □

The algorithm produces a partition $A_1, \ldots, A_k$ of S such that each A_i has maximum density in the complement of $\cup_{j<i} A_j$. The resulting search strategy π orders each A_i in an undetermined manner. The search strategy π is fully determined only if each A_i is a singleton. This only happens in very specific cases, for instance, if f is modular. The maximum density first algorithm then produces the optimal search strategy. It corresponds to Smith's rule [33] for optimal scheduling. In the context of search theory, it corresponds to a result of Blackwell, reported in [21], on the optimal search for an object hidden according to a known probability distribution in one of a finite set of boxes with search costs.

2.2 A better approximation for cost functions of low curvature

We can show that the algorithm of Subsection 2.1 performs better when the cost function f has low *total curvature* (see [36], for example). For $B \subset S$ the *marginal function* f_B is defined as $f_B(A) = f(A \cup B) - f(B)$. The *total curvature* κ of f is

$$\kappa = 1 - \min_{s \in S} \frac{f_{S\setminus\{s\}}(s)}{f(s)}$$

Note that $0 \leq \kappa \leq 1$ and f is modular if and only if the total curvature is 0. For a submodular function, the marginal $f_X(s)$ decreases as $X \subset S$ increases, but it always exceeds $(1-\kappa)f(s)$.

Lemma 4. *Suppose S has maximum search density and f has total curvature κ. Then for all $A \subset S$, the density $\rho(A)$ satisfies*

$$1 - \kappa \leq \rho(A) \leq 1.$$

Proof. The second inequality follows from the fact that S has maximum density of 1. To prove the first inequality, first note that

$$1 = f(\bar{A} \cup A) = f(\bar{A}) + f_{\bar{A}}(A) \geq f(\bar{A}) + (1 - \kappa)f(A) \text{ (by definition of } \kappa).$$

Therefore $(1-\kappa)f(A) \leq 1 - f(\bar{A}) \leq 1 - \mathbf{x}(\bar{A}) = \mathbf{x}(A)$, since $\rho(\bar{A}) \leq 1$, and the inequality follows. □

We can now revisit the proof of Lemma 3 to see that if S has maximum density then for the search π with $S_j = S_j^\pi = \{1, \ldots, j\}$, we have $\rho(S_j) \geq 1 - \kappa$, so

$$c(\pi, \mathbf{x}) = \sum_j \mathbf{x}(j) \frac{\mathbf{x}(S_j)}{\rho(S_j)} \leq \frac{1}{1-\kappa} \sum_j \mathbf{x}(j)\mathbf{x}(S_j) = \frac{1}{1-\kappa} \left(\frac{1}{2} + \frac{1}{2} \sum_j \mathbf{x}(j)^2 \right),$$

which is independent of π. Comparing with (2), we see that the ratio of the expected cost of any search to the expected cost of the optimal search is no more than $1/(1-\kappa)$. Hence, a similar induction argument to Theorem 3 implies the following.

Theorem 4. *Suppose that the submodular function f is given by a value giving oracle, and f has total curvature κ. Then there is a $\left(\frac{1}{1-\kappa}\right)$-approximation for the optimal search strategy against $\mathbf{x}$ that can be computed in strongly polynomial time.*

2.3 An optimal search for fully reducible cost functions

In this section we show how Theorem 1 may be used to determine an optimal search for certain cost functions f that we call *fully reducible*. The idea of fully reducible cost functions is motivated by the following example considered in [3]. Let S be the vertex set of a tree $T = (S, E)$ with edge set E and each $e \in E$ has weight $w(e)$. Let $r \in S$ be the root of the tree. For a set of edges A, define $f(A)$ to be the sum of the edge weights in the tree that is spanned by $\{r\} \cup A$. It is clear that if r has degree 1 then there is an optimal search that begins with the edge incident to r. We generalize this principal by defining *initial sets* below. If r has degree greater than 1, then T is the union of two disjoint subtrees with root r, and it is easy to show that there is a maximal density subset of S whose elements are the vertices of one of these subtrees. So the problem of finding the optimal search can be reduced. We generalize this principal using the concept of separators.

For a set $A \subset S$, we define the *closure* $cl(A)$ of A as the maximal set B containing A such that $f(B) = f(A)$; there is a unique such set. We say that a proper subset $I \subset S$ is an *initial set* if $I \subset cl(s)$ for every $s \in \bar{I}$. We leave it to the reader to check that a set I is an initial set if and only if for any subset $A \subset S$ that contains some element of $\bar{I}$, we have $f(A \cup I) = f(A)$.

Note that if there is an initial set, then the total curvature of f is 1 (that is, the worst possible). Therefore our approximation given in Theorem 4 is not helpful. However, it is easy to show that there is an optimal search with initial segment I.

Lemma 5. *If I is an initial set, then there exists an optimal search with initial segment I.*

Proof. Suppose that π has search order $I_1, A_1, I_2, A_2, \ldots, I_k, A_k$ with $I_i \subset I$ and $A_i \subset \bar{I}$ for all i, and I_1 may be the empty set. Define the new search π' by permuting the order of the I_i and A_i to $I_1, I_2, \ldots, I_k, A_1, A_2, \ldots, A_k$. If $i \in I$, then $S_i^{\pi'} \subset S_i^{\pi}$ so $c(\pi', i) \leq c(\pi, i)$. If $i \in \bar{I}$ then $S_i^{\pi} \cap \bar{I} = S_i^{\pi'} \cap \bar{I}$, so since I is an initial set, $f(S_i^{\pi}) = f(S_i^{\pi'})$ and $c(\pi', i) = c(\pi, i)$. □

Therefore, if an initial set I exists, then it follows from Lemma 2 that in order to find an optimal search of S it is sufficient to find an optimal search of I with respect to f^I and an optimal search of $\bar{I}$ with respect to f_I. This is one way in which f can be reduced.

Finding an initial set can be performed in polynomial time, as we now explain. For any $s \in S$, let I_s be the largest initial set not containing s (if no such initial set exists, let $I_s = \emptyset$). We find I_s; if $I_s = \emptyset$ we proceed to the next element of S, and if all I_s's are empty then there is no initial set.

In order to find I_s, we maintain a candidate T, starting with $T = cl(s)\backslash\{s\}$. We know that $I_s \subseteq T$, and that for every $t \in S\backslash T$, we have $I_s \subseteq cl(t)$. Hence we take an arbitrary $t \in S\backslash T$ and update T as $T \cap cl(t)$.

We iterate this process: while there exists a $t \in S\backslash T$ that we have not yet examined, we update T as $T \cap cl(t)$. We examine every t at most once. At termination, if $T \neq \emptyset$ then we must have $T \subseteq cl(t)$ for every $t \in S\backslash T$, showing that T is an initial set. It is also clear from the construction that $T = I_s$, the largest initial set disjoint from s.

The other way that f can be reduced is by finding a separator, as we now explain. For a lattice $\mathcal{L}$ and $C \subset S$ the restriction to B is $\mathcal{L}^B = \{A \in \mathcal{L} \colon A \subset B\}$. If S and T are disjoint subsets and if $\mathcal{L}$ is a lattice in S and $\mathcal{N}$ is a lattice in T, then the direct sum of these lattices is $\{A \cup A' \colon A \in \mathcal{L},\ A' \subset \mathcal{N}\}$. It is a lattice in $S \cup T$.

Lemma 6. *If B is a separator, then $\mathcal{M}$ is the direct sum of $\mathcal{M}^B$ and $\mathcal{M}^{\bar{B}}$.*

Proof. If A has maximum density, then the inequality

$$\rho^* = \rho(A) = \frac{\mathbf{x}(A \cap B) + \mathbf{x}(A \cap \bar{B})}{f(A \cap B) + f(A \cap \bar{B})} \leq \max\{\rho(A \cap B), \rho(A \cap \bar{B})\}$$

is in fact an equality. It follows that $\rho(A \cap B) = \rho(A \cap \bar{B}) = \rho^*$. □

It follows from Lemma 6 that if B is a separator then either B or $\bar{B}$ (or both) must contain a subset A of density ρ^*. So by Theorem 1, there exists an optimal search of S with initial segment A, where A is a proper subset of S. In that case, we can again apply Lemma 2 to reduce the problem of finding an optimal search into two subproblems.

If a separator exists, then it is possible to find one in strongly polynomial time, using the following method. The *connectivity function* is defined as $d(B) = f(B) + f(\bar{B}) - f(S)$. It is a symmetric non-negative submodular function [11]. We say that a non-empty subset $B \subset S$ is a *split* if $d(B)$ is minimal and $d(A) > d(B)$ for all non-empty $A \subset B$. Obviously, a split is a separator if and only if $d(B) = 0$, and since a split can be computed from a submodular

function minimization, this can be carried out in strongly polynomial time by Schrijver's algorithm, or the Iwata-Fleischer-Fujishige algorithm [14, 28].

We can now define reducibility.

Definition 2. *We say a submodular function f is* reducible *if f has an initial set I or a separator B. We say that f is* fully reducible *if it can be reduced until all B's and I's are singletons.*

We have shown that if f is fully reducible, then by reducing it we can determine an optimal search strategy. We summarize this result with a theorem.

Theorem 5. *Suppose that a submodular function f is given by a value giving oracle. Then we can decide in strongly polynomial time whether f is fully reducible, and if so we can find an optimal search.*

The modular functions are fully reducible by direct sum decomposition, as is the example from [3] described at the beginning of this subsection. Theorem 5 also has an interpretation in scheduling, with regards to scheduling jobs whose precedence constraints are given by a series parallel graph. This is explained in more detail in Subsection 2.4.

2.4 Applications to Scheduling

In a single machine scheduling problem, a set of jobs S has to be ordered in such a way that some cost function is minimized. Every job $s \in S$ has a processing time $p(s) > 0$ and a priority weight $w_j \geq 0$, which we may normalize without loss of generality so that $\sum_{j \in S} w_j = 1$. A *schedule* is an ordering π of the jobs and the completion time of job j is defined as $C_j = \sum_{\pi(i) \leq \pi(j)} p(i)$. If the cost function is $\sum_j w_j C_j$ then an optimal schedule is equivalent to an optimal search for the modular function $f(A) = \sum_{s \in A} p(s)$ and we are back at Smith's rule. However, if the cost function is $\sum_j w_j g(C_j)$ for some non-decreasing function g then nothing is known in general. If $g(C_j) = C_j^\beta$ for $\beta \neq 1$, then we have the scheduling problem $1 \parallel \sum_j w_j C_j^\beta$, considered in [8]. It is unknown whether there exists a polynomial time algorithm to compute an optimal schedule for this problem, or if it is NP hard (see [8] and the references therein).

A scheduling problem may be further complicated by imposing precedence constraints between the jobs, that are specified by a partial order $\prec$. If $s \prec t$ then job t cannot be processed before s is completed. The completion time C_j of job j is defined as before, but now only feasible orderings are permitted. If the cost function is $\sum_j w_j C_j$, then an optimal schedule with precedence constraints is equivalent to an optimal search of S with respect to the submodular cost function $f(A) = p(\tilde{A})$ where $\tilde{A}$ is the precedence closure of A (that is, the set of all jobs that precede some job in A). This scheduling problem $1 \mid prec \mid \sum_j w_j C_j$ is known to be NP hard [17, 19] and it is anticipated that a 2-approximation is best possible [5].

One usually depicts the partial order on the jobs by a Hasse diagram, which is a directed acyclic graph Γ with vertex set S and edges (s,t) if $s \prec t$ and s is an immediate predecessor of t. If Γ is a generalized series parallel graph, then it is possible to compute an optimal schedule in polynomial time [32]. Such a graph is defined recursively by parallel composition and series composition, which corresponds to adding initial sets to a submodular function or taking a direct sum. Our Theorem 3 corresponds to the 2-approximation of a scheduling problem for a general Γ, and Theorem 5 corresponds to a generalized series parallel graph. Thus our results for submodular search overlap known facts from scheduling theory, but we can also add something new:

Theorem 6. *If g is non-decreasing and concave then $1 \mid prec \mid \sum w_j g(C_j)$ admits a 2-approximation that is computable in strongly polynomial time.*

Proof. This is an immediate consequence of Theorem 3, due to the fact that the composition of the concave function g with the submodular function $f : A \mapsto p(\tilde{A})$ is itself submodular. □

It has been brought to the attention of the authors that Theorem 6 has been proved simultaneously in [30] by a different method.

We may be able to do better than a 2-approximation if g defines a submodular cost function f with low total curvature. By way of an example, we take the standard log utility function [26], scaled to fit our framework:

$$g(y) = \frac{log(1+ay)}{log(1+a)},$$

for some positive constant a. Consider the scheduling problem $1 \mid\mid \sum w_j g(C_j)$ with no precedence constraints. This translates into the submodular search problem with a cost function of total curvature $a/(1+a)$, so by Theorem 4, we obtain a $(1+a)$-approximation for the optimal schedule.

3 The Submodular Search Game

We now turn to the submodular search game and we seek optimal mixed strategies for the players. A mixed strategy for the Searcher is some $\mathbf{p}$ which assigns a probabiltiy $\mathbf{p}(\pi)$ to each pure strategy π. We extend the notation for the search cost so that $c(\mathbf{p}, \mathbf{x})$ denotes the expected search cost for mixed strategies $\mathbf{p}$ and $\mathbf{x}$.

Recall that the *base polyhedron* $\mathbb{B}(f)$ of a submodular function $f : S \to \mathbb{R}$ is defined as

$$\mathbb{B}(f) = \{\mathbf{y} \in \mathbb{R}^S : \mathbf{y} \in \mathbb{P}(f), \mathbf{y}(S) = f(S)\}.$$

The base polyhedron is a subset of Hider mixed strategies in the submodular search game. In general, the base polyhedron of a non-negative submodular function is non-empty ([14, page 59]. We apply Theorem 1 to settle a question from [13].

Theorem 7. *Every optimal strategy for the Hider in the submodular search game is in the base polyhedron.*

Proof. By contradiction. Suppose that $\mathbf{x}$ is an optimal Hider strategy, but it is not in the base polyhedron. Then there exists a subset A such that $\mathbf{x}(A) > f(A)$, so that $\rho(A) > 1 = \rho(S)$. It follows that the largest subset M of maximum density is a proper subset of S. Any pure strategy best response π to $\mathbf{x}$ searches M first, by Theorem 1, and hence an optimal mixed strategy of the Searcher assigns positive probability only to orderings of S with intial segment M. Now informally, an optimal response to these Searcher strategies is to hide in $\bar{M}$, which cannot be an equilibrium strategy.

More formally, we observe that every pure search strategy in the support of the optimal strategy $\mathbf{p}$ is a best response to $\mathbf{x}$, so it must start by searching the whole of M. Let $k \in \bar{M}$, then we must have $f(M \cup \{k\}) > f(M)$, otherwise M cannot be maximal. Define $\mathbf{y}$ by $\mathbf{y}(M) = 0, \mathbf{y}(k) = \mathbf{x}(k) + \mathbf{x}(M)$ and $\mathbf{y}(j) = \mathbf{x}(j)$ for any $j \notin M \cup \{k\}$. Then we have

$$c(\mathbf{p}, \mathbf{y}) - c(\mathbf{p}, \mathbf{x}) \geq \mathbf{x}(M)(f(M \cup \{k\}) - f(M)) > 0,$$

so $\mathbf{x}$ cannot be a best response to $\mathbf{p}$: a contradiction. Hence we must have $M = S$ and $\mathbf{x}$ is in the base polyhedron. □

We can also find strategies that are good approximations of the optimal strategies for cost functions with low total curvature. Define the modular function $g(A) = \sum_{s \in A} f(s)$. Then $f(A) \leq g(A)$, and $f(A) \geq (1-\kappa)g(A)$ for all $A \subset S$. It follows that for any mixed strategies $\mathbf{p}$ and $\mathbf{x}$, we have $c_g(\mathbf{p}, \mathbf{x}) \geq c_f(\mathbf{p}, \mathbf{x}) \geq (1-\kappa)c_g(\mathbf{p}, \mathbf{x})$.

The solution to the search game with modular cost function g can be found in [3] or [20]. An optimal strategy for the Hider is given by $\mathbf{x}^*(s) = g(s)/g(S) = f(s)/g(S)$; an optimal strategy $\mathbf{p}^*$ for the Searcher is to begin with an element s with probability $\mathbf{x}^*(s)$ and to search the remaining elements in a uniformly random order.

Proposition 1. *Let κ be the total curvature of f. Then the optimal strategies $\mathbf{x}^*$ and $\mathbf{p}^*$ in the submodular search game with cost function g are $\left(\frac{1}{1-\kappa}\right)$-approximations for the optimal strategies in the submodular search game with cost function f.*

Proof. Let V_f and V_g be the value of the game with cost function f and the game with cost function g, respectively. Then for any strategies $\mathbf{x}$ and p, we have

$$V_g = \max_{\mathbf{x}} c_g(\mathbf{p}^*, \mathbf{x}) \geq \max_{\mathbf{x}} c_f(\mathbf{p}^*, \mathbf{x}) \geq V_f \geq \min_{\mathbf{p}} c_f(\mathbf{p}, \mathbf{x}^*) \geq (1-\kappa) \min_{\mathbf{p}} c_g(\mathbf{p}, \mathbf{x}^*) = (1-\kappa)V_g.$$

It follows that $\max_{\mathbf{x}} c_f(\mathbf{p}^*, \mathbf{x}) \leq V_f/(1-\kappa)$ and $\min_{\mathbf{p}} c_f(\mathbf{p}, \mathbf{x}^*) \geq (1-\kappa)V_f$. □

If f is fully reducible, as defined in Definition 2, we can do much better.

Theorem 8. *Suppose that f is fully reducible. Then the optimal Hider strategy is unique. If f is given by a value giving oracle, then this strategy, $\mathbf{x}^*$ can be computed in strongly polynomial time. An optimal Searcher strategy $\mathbf{p}^*$ can also be computed in strongly polynomial time. The value $V = V_f(S)$ of the game is*

$$V = \frac{1}{2}(f(S) + \Phi), \tag{3}$$

where $\Phi = \Phi_f(S) = \sum_{s \in S} \mathbf{x}^*(s) f(s)$.

Proof. We prove the Hider strategy is unique by induction. If f has an initial set I then the optimal Searcher strategy starts the search in I. Therefore, hiding in $\bar{I}$ dominates hiding in I. To determine the optimal Hider strategy $\mathbf{x}^*$, we can restrict our attention to the game on $\bar{I}$ with cost function f_I, which has a unique Hider strategy that can be calculated in strongly polynomial time, by induction.

If f has no initial set then it must have a separator A, and f can be expressed as a direct sum $f = f^A \oplus f^{\bar{A}}$. We decompose f into $f = f^{A_1} \oplus f^{A_2} \oplus \ldots \oplus f^{A_k}$ such that none of the constituents f^{A_i} decomposes any further. Since $\mathbf{x}$ is in the base polyhedron, S has maximum density, and by Lemma 6 each A_i has maximum density, which uniquely determines $\mathbf{x}^*(A_i) = f(A_i)/f(S)$ for each i. By Lemma 1, none of the A_i contains a proper subset of maximum search density. Therefore, by Theorem 1 an initial segment of an optimal search against $\mathbf{x}^*$ is equal to one of the A_i. In other words, if the Searcher starts a search in A_i, then she continues to search there. Therefore, an optimal response of the Hider is completely specified by the optimal strategies for the subgames on each A_i with cost functions f^{A_i}, and by induction these optimal strategies are unique and can be calculated in strongly polynomial time.

Now we turn to the Searcher's strategy $\mathbf{p}^*$, which we describe recursively. If f has an initial set I, then the Searcher starts by searching I. If f has a separator A, then she searches A first with probability q and otherwise searches $B = \bar{A}$ first, where

$$q = \frac{1}{2} + \frac{\Phi_{f^A}(A) - \Phi_{f^B}(B)}{2f(S)}.$$

We prove by induction that this strategy ensures an expected search cost of at most V, where V is given by Equation (3). First suppose f has an initial set I. Then any best response of the Hider must be an element of $\bar{I}$, so by induction, the expected search cost $C(\mathbf{p}^*, s)$ against an element s of $\bar{I}$ satisfies

$$\begin{aligned} C(\mathbf{p}^*, s) &\leq f(I) + V_{f_I}(\bar{I}) \\ &= f(I) + \frac{1}{2}\left(f_I(\bar{I}) + \Phi_{f_I}(\bar{I})\right) \\ &= f(I) + \frac{1}{2}\left((f(S) - f(I)) + (\Phi_f(S) - f(I))\right) \\ &= \frac{1}{2}(f(S) + \Phi) = V, \end{aligned}$$

Now suppose f has a separator A, and without loss of generality, let $s \in A$. Then, by induction, the expected search cost $C(\mathbf{p}^*, s)$ satisfies

$$\begin{aligned} C(\mathbf{p}^*, s) &\leq V_{f^A}(A) + (1-q)f(B) \\ &= \frac{1}{2}\left(f(A) + \Phi_{f^A}(A)\right) + \left(\frac{1}{2} + \frac{\Phi_{f^B}(B) - \Phi_{f^A}(A)}{2f(S)}\right) f(B) \\ &= \frac{1}{2}\left(f(S) + \mathbf{x}^*(A)\Phi_{f^A}(A) + \mathbf{x}^*(B)\Phi_{f^B}(B)\right) \\ &= \frac{1}{2}(f(S) + \Phi) = V. \end{aligned}$$

This shows that the value is at most V. A similar inductive argument shows that the Hider strategy $\mathbf{x}^*$ ensures the expected search cost is at least V. □

Theorem 8 generalizes results in [3] on expanding search on a rooted tree, where it is shown that the optimal Hider strategy is unique and can be computed efficiently, as can an optimal Searcher strategy. Expanding search on a tree is a submodular search game with a fully reducible cost function.

We may consider the submodular search game in the context of scheduling jobs with processing times and precedence constraints. One player chooses a schedule and the other player chooses a job; the payoff is the completion time of the chosen job. We can interpret this as a robust approach to scheduling, in which one job, unknown to the scheduler, has particular importance, and the scheduler seeks a randomized schedule that minimizes the expected completion time of that job in the worst case. This has a natural application to planning a research project or an innovation process, in which there are many directions the project can take, but it is unknown which task will be fruitful. We call this the *scheduling game*. Theorem 8 gives the solution of the game on series parallel graphs. An interesting direction for future research would be to study the scheduling game on more general partial orders.

4 Final remarks

We have shown that the notion of reducibility is useful for solving both the submodular search problem and the submodular search game. A direction for future research could be to find some measure that captures the "distance" from reducibility, and show that better approximations can be found for cost functions that are close to reducible. It is shown in [6] that better approximations to the single machine scheduling problem $1|prec|\sum w_j C_j$ can be found when the precedence constraints have low *fractional dimension*. It would be interesting to see whether this idea could be generalized to our setting.

The submodular search game that we have studied in this paper is a zero-sum game between one Searcher and one Hider. In search games on networks, one usually restricts attention to one Searcher only, since more Searchers can divide up the space efficiently [1, p 15].

However, in a submodular search game, such a division is impossible and it is interesting to study games with multiple Searchers, which should relate to multi machine scheduling problems. Another extension would be to consider search games with selfish Hiders. Selfish loading games have been studied, and an overview can be found in [35]. These are games between one Searcher and multiple Hiders and a modular payoff function, similar to the scheduling problem $1 \parallel \sum w_j C_j$. A study of submodular search games with selfish Hiders would extend this to $1 \mid prec \mid \sum w_j g(C_j)$, for g concave.

We end with a question. It is known that the complexity of determining an optimal Searcher strategy in a specific search game on a network is NP hard [34], see also [23]. What is the complexity of determining the optimal strategies in the submodular search game?

Acknowledgments.

The authors would like to thank Christoph Dürr for pointing out the connection between expanding search and scheduling.

References

[1] Alpern S, Gal S (2003) The Theory of Search games and Rendezvous. *Kluwer International Series in Operations Research and Management Sciences* (Kluwer, Boston).

[2] Alpern S, Howard JV (2000) Alternating search at two locations. *Dynamics and Control* 10(4):319–339.

[3] Alpern S, Lidbetter T (2013) Mining coal or finding terrorists: the expanding search paradigm. *Oper. Res.* 61(2):265–279.

[4] Alpern S, Lidbetter T (2014) Searching a Variable Speed Network. *Math. Oper. Res.* 39(3):697–711.

[5] Ambühl C, Mastrolili M (2009) Single machine precedence constrained scheduling is a vertex cover problem. *Algorithmica* 53:488-503.

[6] Ambühl C, Mastrolilli M, Mutsanas N, Svensson O (2011) On the approximability of single-machine scheduling with precedence constraints. *Math. Oper. Res.* 36(4):653–669.

[7] Ambühl C, Mastrolili M., Svensson O (2007) Inapproximability results for sparsest cut, optimal linear arrangement, and precedence constrained scheduling. In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*, 329-337.

[8] Bansal N, Dürr C, Thang NK, Vásquez ÓC (2016) The local-global conjecture for scheduling with non-linear cost. *Journal of Scheduling* doi:10.1007/s10951-015-0466-5.

[9] Chekuri C, Motwani R (1999) Precedence constrained scheduling to minimize sum of weighted completion times on a single machine. *Discr. Appl. Math.* 98(1):29–38.

[10] Chudak FA, Hochbaum DS (1999) A half-integral linear programming relaxation for scheduling precedence-constrained jobs on a single machine. *Oper. Res. Lett.* 25:199–204.

[11] Cunningham W (1983) Decomposition of submodular functions. *Combinatorica* 3(1):53–68.

[12] Dürr C, Jèz Ł, Vásquez ÓC (2015) Scheduling under dynamic speed-scaling for minimizing weighted completion time and energy consumption. *Discrete Appl. Math.* 196:20–27.

[13] Fokkink R, Ramsey D, Kikuta K (2016) The search value of a set. *Annals Oper. Res.* 1–11.

[14] Fujishige S (2005) Submodular functions and Optimization. *Annals Discr. Math.* 58, Elsevier.

[15] Gal S (1980) Search games (Academic Press, New York).

[16] Gal S (2011) Search games. *Wiley Encylopedia of OR and MS* (John Wiley & Sons, Hoboken, NJ).

[17] Garey MR, Johnson DS (1979) Computers and Intractability: A Guide to the Theory of NP-completeness (Freeman, San Francisco).

[18] Hall LA, Schulz AS, Shmoys DB, Wein J (1997) Scheduling to minimize average completion time: off-line and on-line algorithms. *Math. Oper. Res.* 22(3):513–544.

[19] Lawler EL, Lenstra JK, Rinnooy Kan AHG, Shmoys DB (1993) Sequencing and scheduling: algorithms and complexity. *Handbooks in Operations Research and Managament Science 4* Elsevier, 445–522.

[20] Lidbetter T (2013) Search Games with Multiple Hidden Objects, *SIAM J. Control and Optim.* 51(4):3056–3074.

[21] Matula D (1964) A periodic optimal search. *Am. Math. Mon.* 71(1):15–21.

[22] Margot F, Queyranne M, Wang Y (2003) Decompositions, network flows and a precedence constrained single machine scheduling problem. *Oper. Res.* 51(6):981–992.

[23] Megiddo N, Hakimi SL, Garey MR, Johnson DS, Papadimitriou CH (1988) The complexity of searching a graph. J. ACM 35(1):18–44.

[24] Megow N, Verschae J (2013) Dual techniques for scheduling on a machine with varying speed. In *Proc. of the 40th International Colloquium on Automata, Languages and Programming (ICALP)*, 745–756.

[25] Nagano K (2007) A strongly polynomial algorithm for line search in submodular polyhedra. *Discr. Optim.* 4(3):349–359.

[26] von Neumann J, Morgenstern O (2007) Theory of Games and Economic Behavior (Princeton University Press).

[27] Pisaruk NN (2003) A fully combinatorial 2-approximation algorithm for precedence-constrained scheduling a single machine to minimize average weighted completion time. *Discrete Appl. Math.* 131(3):655–663.

[28] Schrijver A (2003) Combinatorial optimization, polyhedra and efficiency, *Algorithms and Combinatorics 24* (Springer, Berlin).

[29] Schulz AS (1996) Scheduling to minimize total weighted completion time: Performance guarantees of LP-based heuristics and lower bounds. In *Proceedings of the 5th Conference on Integer Programming and Combinatorial Optimization (IPCO)*, 301–315.

[30] Schulz AS, Verschae J (2016) Min-sum scheduling under precedence constraints. In *European Symposia on Algorithms, 2016 (ESA 2016)*, forthcoming.

[31] Schuurman P, Woeginger GJ (1999) Polynomial time approximation algorithms for machine scheduling: ten open problems. *Journal of Scheduling* 2(5):203–213.

[32] Sidney J (1975) Decomposition algorithms for single-machine sequencing with precedence relations and deferral costs. *Oper. Res.* 23(2):283–298.

[33] Smith WE (1956) Various optimizers for singlestage production. *Naval Res. Logist. Quarterly* 3(1-2):59–66.

[34] von Stengel B, Werchner R (1997) Complexity of searching an immobile Hider in a graph. *Discrete Appl. Math.* 78(1):235–249.

[35] Vöcking B (2007) Selfish load balancing. In *Algorithmic Game Theory*, eds. Nisam N, Roughgarden T, Tardos É, Vazirani VV (Cambridge University Press), 517–542.

[36] Vondrák J (2010) Submodularity and curvature: the optimal algorithm. *RIMS Kokyuroku Bessatsu B* 23:253–266.